Machine Learning aids rapid assessment of aftershocks: Application to the 2022-2023 Peace River earthquake sequence, Alberta, Canada

Jinji Li, Jesús Rojas-Parra, Rebecca O. Salvage, David W. Eaton, Kris A. Innanen, Yu Jeffrey Gu, and Wenhan Sun

ABSTRACT

The integration of machine learning (ML) models has ignited a paradigm shift in seismic analysis, fostering enhanced efficiency in capturing patterns of seismic activity with reduced need for time-consuming user interaction. Here, we investigate automated event detection and extraction of seismic phases using two widely used ML models, EQTransformer and PhaseNet. We applied both models to four weeks of continuous recordings of aftershocks using a temporary array following the November 30, 2022 M_L 5.6 earthquake near Peace River, Alberta, Canada. Both tools identified $>$1000 events over the recording period. The aftershocks are located in close proximity and depth to the M_L 5.6 mainshock on November 30, 2023, as well as to disposal operations that were ongoing at the time. Although there are some differences in the temporal and spatial evolution of the detected events by each model, both sets of detections reveal similar patterns of the aftershock distribution that were not identified by the regional network. Our results highlight the advantages of using ML models for rapid detection and assessment of seismicity following felt events, which is important for assessing hazard potential and risk in near real time.

INTRODUCTION

Detecting events and picking seismic phases are fundamental to seismological workflows (Bormann, 2012). Accurately identifying and characterizing seismic events and precise phase picking are critical to discern earthquake dynamics, subsurface structures, and seismic activity patterns (Mousavi et al., 2020). As such, the integration of machine learning (ML) models for seismic event detection and phase picking represents a paradigm shift in seismology, offering a promising alternative to traditional approaches. For example, Karimzadeh et al. (2019) used multiple ML algorithms, including naive Bayes, K-nearest neighbors, support vector machine, and random forests, together with the information of slip distribution, active faults' locations, and Coulomb stress to predict earthquakes' aftershock patterns. The results indicate that the aftershock pattern prediction is possible even with a small database. Zhu and Beroza (2018) developed a deep-neural-network-based arrival picking model named PhaseNet, which analyzes three-component seismic waveforms to provide accurate arrival times of P and S waves. Woollam et al. (2019) presented a convolutional neural network (CNN) for classifying seismic phase onsets for local seismic networks. With a rather small training dataset, this CNN-based approach outperforms one of the classical approaches. Mousavi et al. (2019a, 2020) introduced two ML-based models, namely the CNN-RNN Earthquake Detector (CRED) and the Earthquake Transformer (EQTransformer). The former was successfully applied to a continuous dataset recorded in Central Arkansas and tested to be efficient and promising in earthquake detection, and the latter has shown the potential for detecting and characterizing more and smaller events. Recent research has shown that by harnessing the power of deep neural networks, ML

methods can efficiently and objectively extract subtle patterns and features from large volumes of seismic data that might elude human observers. This degree of automation significantly reduces routine workload and enhances the precision and consistency of seismic event detection and phase picking, thus contributing to more robust earthquake catalogs and improving seismicity analysis and seismic hazard assessments (Banna et al., 2020). Additionally, the ability of ML models to analyze large-scale and continuous seismic data streams in real-time makes them invaluable tools for characterizing geological susceptibility and monitoring induced seismicity (Wozniakowska and Eaton, 2020; Limbeck et al., 2021; Fasola and Brudzinski, 2023).

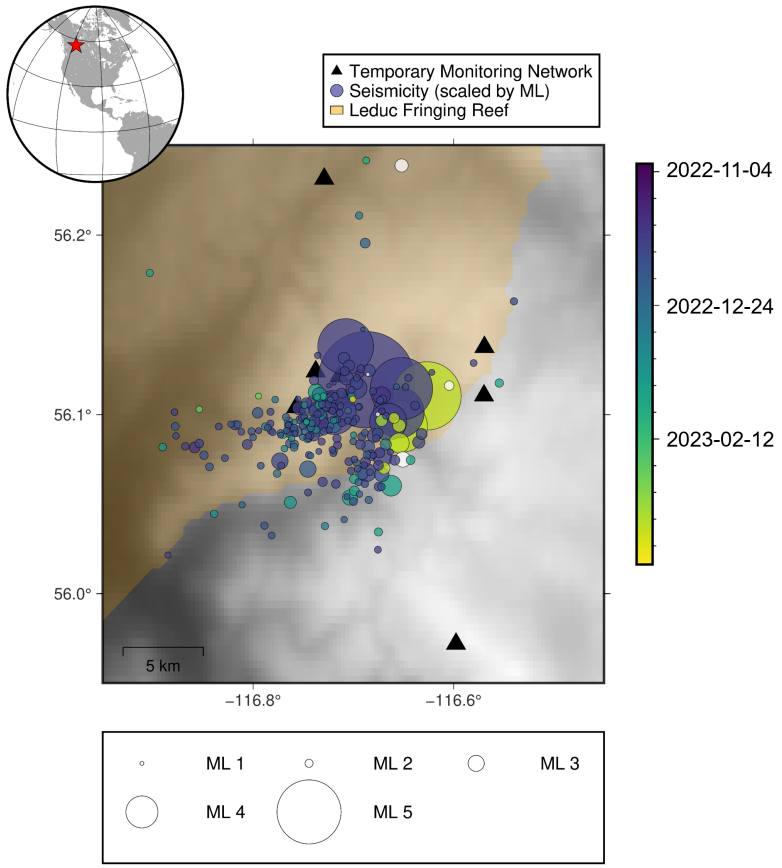

FIG. 1. Seismicity (November 1, 2022 to May 1, 2023) associated with the November 30, 2022 M_L 5.6 Peace River earthquake sequence (Schultz et al., 2023; Salvage et al., 2023; Vasyura-Bathke et al., 2023). Seismicity was determined by a standard STA/LTA algorithm on regional seismograph stations and is coloured by time and scaled by local magnitude, allowing a direct comparison to ML methods (Figure 2). The extent of the Leduc fringing reef, a disposal zone for saltwater injection that surrounds the Peace River Arch, is shown in orange (Alberta Geological Survey, 2019). Temporary monitoring stations deployed by the Alberta Geological Survey ∼1 week after the mainshock are shown as black triangles. *Inset*: Location of the study area (red star).

In this paper, we evaluate the effectiveness of ML models when applied to a temporary array deployed to record aftershocks. The data were acquired by the Alberta Energy Reg-

ulator (AER) in the Peace River (PR) region of north-central Alberta, Canada (Figure 1), using eight portable three-component broadband stations (sampling frequency of 250 Hz) deployed rapidly in the aftermath of a M_L 5.6 earthquake on November 30, 2022 (Schultz et al., 2023; Vasyura-Bathke et al., 2023; Salvage et al., 2023). Stations were deployed within 15 km of the mainshock location. Continuous seismic data from December 7, 2022 to January 13, 2023, inclusive, used in this study was generously provided by the Alberta Geologic Survey. This event (sequence) represents one of the largest recorded seismic events in Alberta and is controversial in nature as it was originally determined to be a natural event (Alberta Energy Regulator, 2022), but subsequent analysis has suggested that it may have been induced by nearby disposal operations (Schultz et al., 2023; Vasyura-Bathke et al., 2023; Salvage et al., 2023). Our aim is to evaluate and compare the effectiveness of two widely used ML-based systems, EQTransformer and PhaseNet, for rapid analysis of aftershock sequences.

EARTHQUAKE DETECTION WITH ML MODELS: EQTRANSFORMER AND PHASENET

In the realm of ML models designed to automate seismic event detection and phase picking, two widely used methods are the Earthquake Transformer (EQTransformer) (Mousavi et al., 2020) and PhaseNet (Zhu and Beroza, 2018). These have emerged as leading tools in the field, offering significant advancements in analysis of seismicity. In a comprehensive study conducted by Munchmeyer et al. (2022), an evaluation of various ML approaches demonstrated that EQTransformer and PhaseNet out-performed most other existing models (e.g. CNN-RNN Earthquake Detector (Mousavi et al., 2019a); Generalized Phase Detection (Ross et al., 2018)) in event detection, phase identification and in the determination of event onset timing.

Drawing inspiration from convolutional and recurrent neural networks (CNNs and RNNs), EQTransformer has demonstrated proficiency in automating seismic event detection and precise phase picking from very large datasets. Its ability to discern intricate temporal relationships within seismic waveforms enables it to detect subtle seismic signals that may prove challenging for conventional event detection methods (Munchmeyer et al., 2022; Mousavi et al., 2020). Similarly, PhaseNet has emerged as a powerful model for phase picking, exhibiting its capability to extract phase arrival times from seismic waveforms. Leveraging U-Net architecture, PhaseNet achieves high performance in accurately identifying P- and S-phases (Zhu and Beroza, 2018). To associate phase picks with events, we used the Gaussian Mixture Model Associator (GaMMA), which assumes a hyperbolic move-out of arrival times of different phases and amplitudes (Zhu et al., 2022).

In the evaluation of both EQTransformer and PhaseNet using the Yangbi and Maduo earthquake datasets, Jiang et al. (2021) found that neural networks with deeper layers and complex structures may not necessarily enhance earthquake detection performance. Hence, utilization and analysis of more local datasets are needed for a more comprehensive understanding and further development of new models. In the following sections, we employ previously trained EQTransformer and PhaseNet models to analyze aftershocks following the M_L 5.6 earthquake on November 30, 2022 in Alberta, followed by a performance comparison.

MODEL IMPLEMENTATION

Hyper-parameter selection significantly impacts model performance, but finding the optimal configuration can vary depending on the specific model and is often computationally demanding (Wu et al., 2019). To ensure meaningful comparisons between models, we adopt fixed model architectures in order to focus solely on model performance by selecting (to the extent possible) identical hyper-parameters. This strategy enables us to concentrate on evaluating and directly comparing the effectiveness of the EQTransformer and PhaseNet models for automated seismicity analysis tasks, providing insights into their respective capabilities. Following this, we adopt the Seisbench interface (Woollam et al., 2022) for its role in standardizing the workflow and granting access to a diverse array of cutting-edge seismological machine-learning models and datasets. The seisbench API contains several pre-trained models, including both EQTransformer and PhaseNet. In this study, both models have undergone pre-training using the well-established STanford EArthquake Dataset (STEAD) (Mousavi et al., 2019b), a dataset acknowledged for its suitability in evaluating these models' performance (Munchmeyer et al., 2022). STEAD contains a substantial dataset, encompassing an impressive 1.2 million traces, including 450,000 earthquakes (ranging from magnitude 0.5 to 8) and approximately 1 million P and S picks.

As a result of pre-testing trials, the hyperparameters employed in this study were chosen to balance the speed and performance. The parameters used in this study are: batch size: 512; overlap: 256 P-threshold: 0.55; S-threshold: 0.55; the number of CPUs: 10. After the evaluation with CPUs, we subsequently incorporated CUDA acceleration into both models. For EQTransformer, we used a detection threshold of 0.7; all other parameters remain identical irrespective of the model architectures.

After undergoing pre-processing steps, including linear de-trending, down-sampling to 100 Hz to fit the pre-trained models, and high-pass filtering excluding signals above 1 Hz, the two machine learning models were applied using the parameters specified above to determine phase picks. We then used the Gaussian Mixture Model Association (GaMMA) (Zhu et al., 2022) for phase association. For this method, each earthquake is represented as a cluster encompassing P and S phases that exhibit an approximately hyperbolic moveout of arrival times and a decline in amplitude as a function of distance. An event's underlying distribution of phase selections is characterized using a multivariate Gaussian distribution, where the mean values are dictated by the anticipated arrival time and amplitude derived from the seismic event responsible for the observations. PhaseNet necessitates the input of an average velocity model. In this study, we have adopted regional velocity averages of 4868.5 m/s for P-wave velocity and 2863.8 m/s for S-wave velocity based on the information in recent research from Schultz et al. (2023); Salvage et al. (2023). To process the data effectively, we apply the DBSCAN (Ester et al., 1996) clustering option within GaMMA. This method excels at grouping data points in close proximity while efficiently identifying noise and outliers. The minimum sample parameter for DBSCAN is set to 3. Additionally, the following hyperparameters were selected for filtering lower quality associations: minimum picks: 6; minimum P-picks: 3; minimum S-picks: 3; maximum values for phase time residual, phase-amplitude residual, and covariance term, namely σ_{11}, σ_{22}, and σ_{12}: 2.0, 1.0, and 1.0, respectively. GaMMA is primarily employed for the initial association process, with the majority of parameters being set to their default values. Details about the

hyperparameters can be found in Zhu and Beroza (2018).

After event detection and phase association using the ML-based models, we used NonLinLoc (Lomax et al., 2000, 2009) to compute hypocenter locations. NonLinLoc is a global-search method that employs a probabilistic framework to determine event location using estimated posterior probability density functions. We use the regional velocity model by Schultz et al. (2023) for hypocenter determination.

RESULTS

Using the approach described above, EQTransformer had 12,578 pickings with all the stations and detected a total of 1,241 events, while PhaseNet had 14,729 pickings with all the stations and detected 1,078 events in total (Figure 2). The average processing time per station with EQTransformer is 7.63 minutes with 10 CPU cores, which is approximately 1.8 times longer than with PhaseNet (4.07 minutes). When utilizing 1 CUDA core, the processing time for EQTransformer and PhaseNet is 2.63 minutes and 2.22 minutes, respectively. Despite these differences, the time and space distributions of the associated aftershock events appear to be comparable. This is illustrated by Figure 2, which shows that despite differences in internal architecture of these methods, the event localization results produced by both models are broadly similar. In particular, PhaseNet and EQTransformer yield NonLinLoc-based results in proximity to the mainshock on November 30, 2022 in both space and depth, as well as within the area of the Leduc fringing reef and close to its boundary, and close to several active disposal wells in the area, which has also been shown to be the case for aftershocks detected by the regional network (Vasyura-Bathke et al., 2023; Salvage et al., 2023).

Despite EQTransformer identifying more events during the $\sim$4 week recording period, errors associated with the hypocenter locations were significantly worse, with an average error in the X, Y direction of 6.2 km and an average error in Z of 5.5 km compared to PhaseNet (average X, Y error of 0.9 km and 1.1 km in Z). This is reflected in the spatial extent of hypocenter locations (Figure 2a,b,d,e), where events identified by PhaseNet appear much more clustered in both area and depth. However, artifacts relating to the implementation of the grid search algorithm are evident in depth, in particular for both PhaseNet (Figure 2b) and EQTransfomer (Figure 2e), as revealed by several unlikely linear features. Temporally, events detected by PhaseNet and EQTransformer (2c and f, respectively) appear moderately concurrent, with the maximum daily event counts occurring between December 10 and December 15, 2022, although EQTransformer did detect events earlier in the recording period (from December 7 onwards, Figure 2f) than PhaseNet (Figure 2c).

To ensure the identification of common events detected by both models, a time tolerance was implemented to determine if two events occurred within close temporal proximity. Specifically, when the start times of two events fall within this tolerance, they are considered a single event. With a tolerance of 1 second, a total of 838 common events were identified, indicating that EQTransformer and PhaseNet individually detected 404 and 240 events, respectively. By applying a relatively broad 1-second time window as the tolerance, the spatial distribution of common events unveils a substantial overlap. For PhaseNet, the hypocenter locations' latitude, longitude, and depth ranges extend from $56.09 - 56.13$ de-

grees, $-116.74 - 116.63$ degrees, and 1.48 km - 5.89 km, respectively. In the case of EQTransformer, the hypocenter locations' latitude, longitude, and depth ranges encompass $56.08 - 56.13$ degrees, $-116.73 - 116.61$ degrees, and 1.44 - 5.4 km, respectively. Reducing the tolerance to 0.5 seconds resulted in 794 common events, with 448 events identified by EQTransformer and 284 by PhaseNet, providing a refined assessment of event overlaps between the two models. The observed overlap, characterized by common events, suggests that EQTransformer in this instance, exhibits greater sensitivity to noisy events, resulting in a more dispersed spatial distribution in Figure 2c and d.

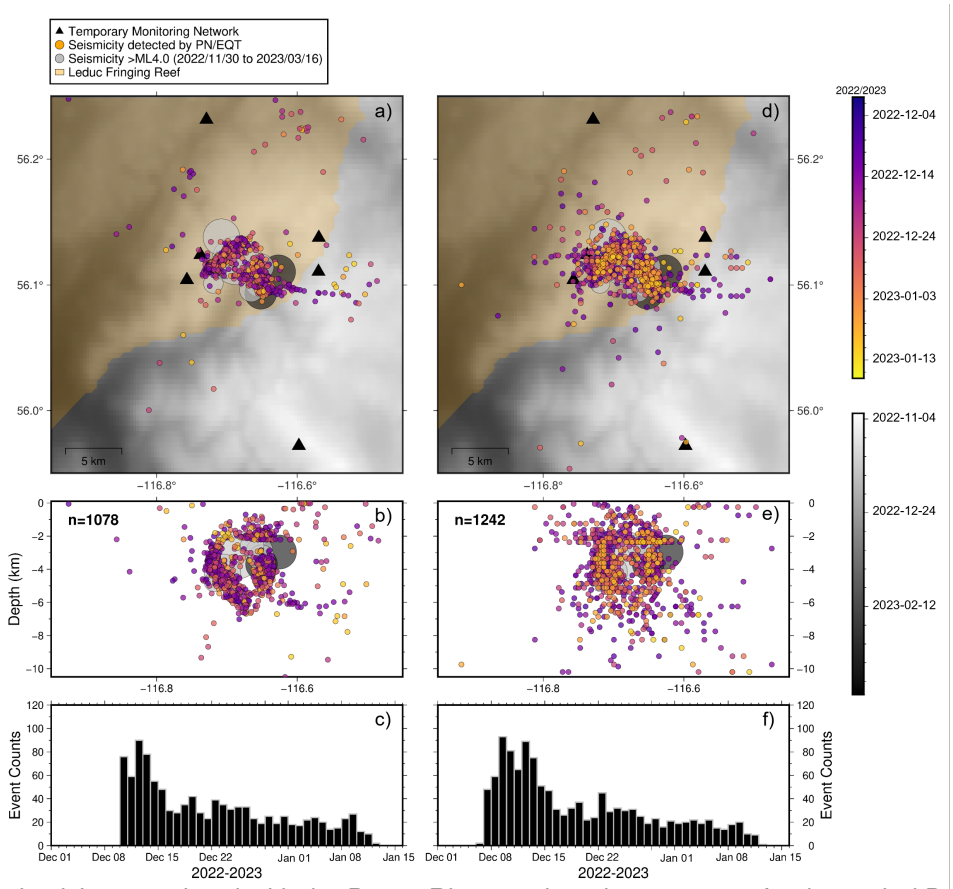

FIG. 2. Seismicity associated with the Peace River earthquake sequence for the period December 7, 2022, to January 13, 2023, determined in this study using data from temporary stations deployed by the Alberta Geological Survey ∼1 week after the mainshock on November 30, 2022. Events are colored by time; portable broadband stations are shown as black triangles, and the extent of the Leduc fringing reef is shown in orange (Alberta Geological Survey, 2019). Seismic events detected using a regional network with $M_L \geq 4.0$ from November 30, 2022 to March 31, 2023 are shown in grey and scaled by magnitude (Vasyura-Bathke et al., 2023). Map and depth view of events detected by PhaseNet (a,b) and EQTransformer (d,e), with 1,078 events detected by PhaseNet and 1,242 events detected by EQTransformer. Histograms of daily event counts for PhaseNet (c) and EQTransformer (f) suggest comparable daily event counts using both methods.

As an illustrative example, Figure 3 provides a segment of the recorded signal from station 3 corresponding to an event concurrently detected by both models. Based on the findings obtained from NonLinLoc, this event appears to have occurred on December 19, 2022 at 05:06:46. Remarkably, station 3 managed to capture this event within a mere one-second timeframe. Although the overall detection results using these methods display a

high degree of similarity, subtle discrepancies do arise, primarily owing to the intricate nature of waveform patterns and the inherent capabilities of these models in capturing aftershock events. From this figure, in general, the PhaseNet model registers slightly delayed arrival times for both P and S arrivals in comparison to EQTransformer, with a larger discrepancy observed in the identification of P arrivals. This systematic bias could stem from various factors, including the selection of hyperparameters in machine learning models, but is primarily attributed to the distinct architectures of the two models. Despite configuring comparable parameter sets, the mathematical formulations might assign varying weights to different factors, contributing to the observed discrepancies. This observation hints at the possibility of requiring human intervention for further evaluation following the initial rapid assessment.

In Figure 4, we present events that were independently detected by EQTransformer and PhaseNet. Notably, these events transpired on distinct dates and were both captured by station 10. The event identified by PhaseNet took place on December 26, 2022, at 18:34:04, with station 10 recording it in slightly less than 1 second. However, the EQTransformer-detected event occurred on December 9, 2022, at 11:24:59, and station 10 recorded it in approximately 0.3 seconds. Thought-provokingly, the waveform clarity suggests that these events should have been detected by both models, emphasizing the need for human intervention following rapid evaluation via machine learning techniques.

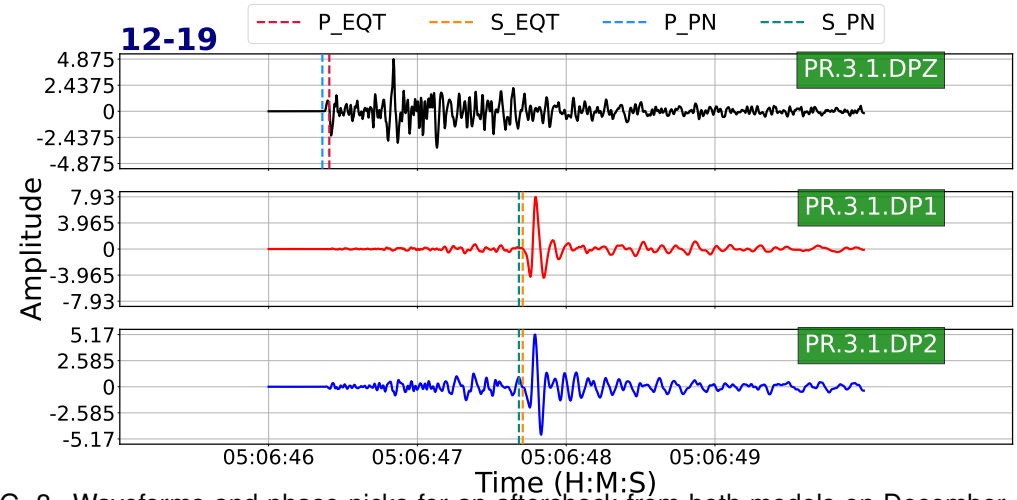

FIG. 3. Waveforms and phase picks for an aftershock from both models on December 19, 2023. The black, red, and blue waveform plots are three different components, as indicated by the upper-right green boxes in each subplot. The dark-red and orange dashed lines are P and S phase identifications by EQTransformer. The light-blue and green dashed lines are P and S phase identifications by PhaseNet. The length of this snippet is around 4 seconds.

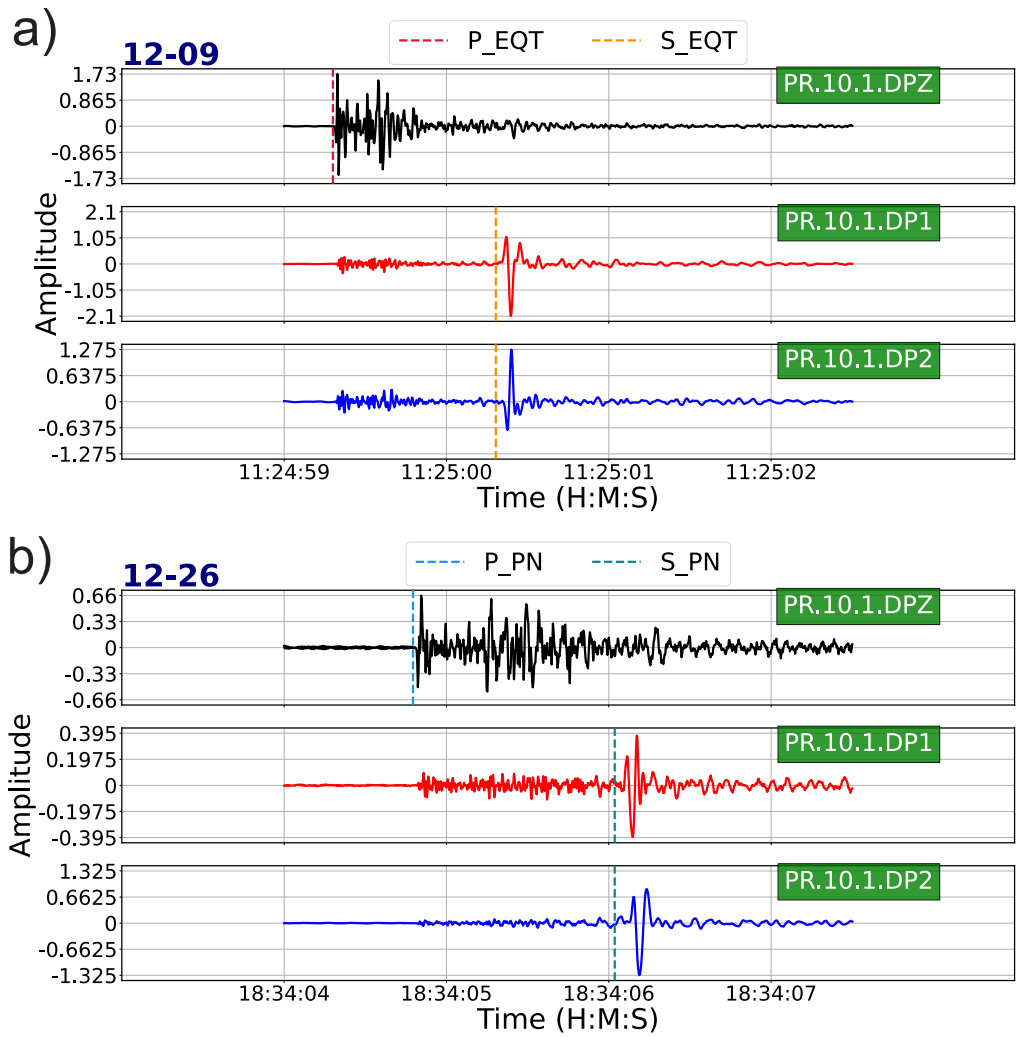

FIG. 4. Waveforms and phase picks for aftershocks independently detected by EQTransformer and PhaseNet on December 9 and December 26, 2023. The black, red, and blue waveform plots are three different components, as indicated by the upper-right green boxes in each subplot. The length of each snippet is around 4 seconds. a) Aftershock event captured by EQTransformer. The dark-red and orange dashed lines are P and S phase identifications by EQTransformer. b) Aftershock event captured by PhaseNet. The light-blue and green dashed lines are P and S phase identifications by PhaseNet.

DISCUSSION

While the total count of associated events from both models appears similar, it's essential to note that these results arise from fundamentally different approaches. EQTransformer operates in a manner that encompasses both event detection and subsequent phase picking, with the event detection preceding the phase picking stage. Consequently, all phase picks are temporally confined to an event detection window. Conversely, PhaseNet's principal role revolves around the autonomous identification of P and S arrivals, decoupled from the event detection process. In terms of uncertainty quantification, EQTransformer typically generates probabilities for event detection and phase picking by the neural network model's output layer and represents the likelihood of the presence of specific events

or phases (Mousavi et al., 2020), while PhaseNet converts the waveforms into probability distributions with several spikes of P and S arrivals (Zhu and Beroza, 2018). Such distinction underscores the inherent differences in the outputs of these two pickers and the nuanced nature of their respective methodologies, thus highlighting the importance of selecting the most suitable model for specific seismic event analysis tasks.

In the context of this study, we have purposefully chosen to maintain consistency by setting the hyperparameter for both models in a comparative way, ensuring a relatively equitable basis for comparison. Based on our current parameter selection for rapid assessment of aftershocks, it's noteworthy to mention that PhaseNet outperforms EQTransformer in this context. PhaseNet offers the advantage of a shorter processing time and demonstrates a reduced error in event location as verified in Figure 2, which we will elaborate on in the next paragraph. However, in practice, optimizing model efficiency can be further enhanced through approaches like grid searches for fine-tuning hyperparameters, which can be the next step of our research. It is also worth mentioning that EQTransformer, while not surpassing PhaseNet in this comparison, does have its merits, particularly in achieving slightly closer phase determinations. However, the determination of which model yields more precise and reasonable phase picks may require manual evaluation to provide definitive justification.

The consistency of results displayed in Figure 2 implies that the seismic events are concentrated close to active disposal wells on the edge of the Leduc fringing reef. When considering event focal depths, it is evident that the predominant cluster of earthquakes aligns with $M_L \geq 4.0$) seismic events detected from the regional network from November 2022 through to March 2023 (Figure 2a,b,d,e; grey circles). Locations obtained for events detected using PhaseNet appear to depict several linear features extending from $\sim$2.5 to $>$6 km at depth (Figure 2b), which may reflect several sub-parallel fault systems that have been identified by Vasyura-Bathke et al. (2023) in this area using aftershocks from the same sequence identified using the regional network. The smaller hypocenter location errors for events detected by PhaseNet allow additional support for the evidence of these structures.

Our findings in this study demonstrate the potential of machine learning for rapid seismic event assessment, particularly when handling extensive datasets. The primary advantage of employing machine learning in this context is the capacity to process large volumes of seismic data with remarkable efficiency. This is especially critical when confronted with datasets of considerable scale, a situation where manual assessment methods are impractical. Machine learning models can quickly analyze and classify seismic events and phase arrivals, saving time and resources. However, it is important to acknowledge that the ML models employed in this study were initially trained on extensive global datasets, which may not be optimally tailored to account for the specific characteristics of local datasets. Local datasets may exhibit distinct characteristics when compared to the training data used for these models, including noise patterns and signal bandwidth. Consequently, these models may not be finely tuned to capture the specific characteristics of localized datasets, which can present unique challenges. These discrepancies can potentially influence the model's performance, and therefore, it's imperative to be cautious when applying such models to local seismic data. Nonetheless, our preliminary findings suggest that these models exhibit a remarkable degree of consistency and efficiency in their performance,

even when applied to local data. This underscores the adaptability and generalization capabilities of these machine learning models. It's also a testament to the robustness of these models when confronted with variations in seismic data characteristics. Nevertheless, to ensure the most accurate and reliable results, it's crucial to recognize and account for the potential distinctions between local and global datasets, particularly in cases where these distinctions can impact the desired outcomes of the analysis. Training models for achieving localized monitoring thus become another possible direction for further study.

CONCLUSIONS

We applied two popular ML models for event detection and phase picking, EQTransformer (Mousavi et al., 2020) and PhaseNet (Zhu and Beroza, 2018), to continuous waveform data from a temporary aftershock deployment following the 2022/11/30 M_L 5.6 Peace River earthquake. Using identical (to the extent possible) choices of hyper-parameters, these two methods yielded a similar number of event detections: EQTransformer detected a total of 1,241 events, while PhaseNet detected 1,078 events. Although the inferred hypocentre distributions reveal some differences when considered in detail, this comparison reveals a consistent overall pattern of aftershocks. Temporally, event detections peak in the first few days after station deployment (December 10-15, 2022). Spatially, hypocenters occur in close proximity to both the edge of the Leduc fringing reef and several active water disposal wells. In depth, events are clustered close to the M_L 5.6 mainshock event and several significant ($> M_L$ 4.0) fore- and aftershocks, detected by regional stations. In general, the application of ML models shows a high degree of promise for near real-time analysis of aftershocks.

DATA AND RESOURCES

Continuous waveform data and related metadata from December 7, 2022 to January 13, 2023 (inclusive) from a temporary seismometer array were generously provided by the Alberta Geological Survey, part of the Alberta Energy Regulator (AER). The data were acquired by the AER in collaboration with the Dr. Yu Jeffrey Gu at the University of Alberta. The authors gratefully acknowledge the contributions of the AER to the acquisition of this dataset, and are deeply grateful that it was provided freely for this study by Chris Filewich (AER). Data processing and analysis were conducted using PhaseNet (Zhu and Beroza, 2018) and Phase Association (Zhu et al., 2022), EQTransformer (Mousavi et al., 2020), Obspy (Beyreuther et al., 2010) and NonLinLoc v7 (Lomax et al., 2000, 2009); all of which are open source. Figures 1 and 2 were produced using PyGMT (Uieda et al., 2023).

ACKNOWLEDGEMENTS

This work was supported by CREWES (Consortium for Research in Elastic Wave Exploration Seismology) and C-DaPS (Consortium for Distributed and Passive Sensing), as well as the Natural Science and Engineering Research Council of Canada (NSERC) through grants CRDPJ 543578-19 and ALLRP-548576-2019.

One of the authors of this report was supported by the CSEG Foundation.

REFERENCES

Alberta Energy Regulator, 2022, Seismic Events Southeast of Peace River, Online, last Accessed: 30 May 2023.
URL https://www.aer.ca/providing-information/news-and-resources/news-and-announcements/announcements/announcement-november-30-2022

Alberta Geological Survey, 2019, 3D Provincial Geological Framework Model of Alberta, version 2, Online, AER/AGS Model 2018-02. Last Accessed: 30 May 2023.
URL https://ags.aer.ca/publication/3d-pgf-model-v2

Banna, M. H. A., Taher, K. A., Kaiser, M. S., Mahmud, M., Rahman, M. S., Hosen, A. S. M. S., and Cho, G. H., 2020, Application of artificial intelligence in predicting earthquakes: State-of-the-art and future challenges: IEEE Access, **8**, 192,880–192,923.

Beyreuther, M., Barsch, R., Krischer, L., Megies, T., Behr, Y., and Wassermann, J., 2010, Obspy: A Python toolbox for seismology: Seismol. Res. Lett., **81**, No. 3, 530–533.

Bormann, P., 2012, New manual of seismological observatory practice (NMSOP-2): GeoForschungsZentrum Potsdam, Potsdam, Germany.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X., 1996, A density-based algorithm for discovering clusters in large spatial databases with noise, *in* Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96, AAAI Press, 226–231.

Fasola, S. L., and Brudzinski, M. R., 2023, Machine learning reveals additional hydraulic fracture-induced seismicity in the Eagle Ford shale: Journal of Geophysical Research: Solid Earth, **128**, No. 2, e2022JB025,436, https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2022JB025436.

Jiang, C., Fang, L., Fan, L., and Li, B., 2021, Comparison of the earthquake detection abilities of PhaseNet and EQTransformer with the Yangbi and Maduo earthquakes: Earthquake Science, **34**, No. 5, 425–435.

Karimzadeh, S., Matsuoka, M., Kuang, J., and Ge, L., 2019, Spatial prediction of aftershocks triggered by a major earthquake: A binary machine learning perspective: ISPRS International Journal of Geo-Information, **8**, No. 10.
URL https://www.mdpi.com/2220-9964/8/10/462

Limbeck, J., Bisdom, K., Lanz, F., Park, T., Barbaro, E., Bourne, S., Kiraly, F., Bierman, S., Harris, C., Nevenzeel, K., Bezemer, T., and van Elk, J., 2021, Using machine learning for model benchmarking and forecasting of depletion-induced seismicity in the Groningen gas field: Computational Geosciences, **25**.

Lomax, A., Michelini, A., and Curtis, A., 2009, Earthquake Location, Direct, Global-Search Methods, Springer New York, New York, NY, 1–33.

Lomax, A., Virieux, J., Volant, P., and Berge-Thierry, C., 2000, Probabilistic Earthquake Location in 3D and Layered Models, Springer Netherlands, Dordrecht, 101–134.
URL https://doi.org/10.1007/978-94-015-9536-0_5

Mousavi, S., Ellsworth, W., Weiqiang, Z., Chuang, L., and Beroza, G., 2020, Earthquake transformer—an attentive deep-learning model for simultaneous earthquake detection and phase picking: Nature Communications, **11**, 3952.

Mousavi, S., Weiqiang, Z., Sheng, Y., and Beroza, G., 2019a, CRED: A Deep Residual Network of Convolutional and Recurrent Units for Earthquake Signal Detection: Scientific Reports, **9**.

Mousavi, S. M., Sheng, Y., Zhu, W., and Beroza, G. C., 2019b, STanford EArthquake Dataset (STEAD): A Global Data Set of Seismic Signals for AI: IEEE Access, **7**, 179,464–179,476.

Munchmeyer, J., Woollam, J., Rietbrock, A., Tilmann, F., Lange, D., Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Michelini, A., Saul, J., and Soto, H., 2022, Which Picker Fits My Data? A Quantitative Evaluation of Deep Learning Based Seismic Pickers: Journal of Geophysical Research: Solid Earth, **127**, No. 1, e2021JB023,499, e2021JB023499 2021JB023499, https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2021JB023499.

Ross, Z. E., Meier, M.-A., Hauksson, E., and Heaton, T. H., 2018, Generalized seismic phase detection with deep learning: Bulletin of the Seismological Society of America, **108**, No. 5A, 2894–2901.

Salvage, R. O., Eaton, D. W., Furlong, C. M., Dettmer, J., and Pedersen, P. K., 2023, Induced or Natural? Towards Rapid Expert Assessment, with Application to the M_W 5.2 Peace River Earthquake Sequence: Seismological Research Letters, in Review.

Schultz, R., Woo, J.-U., Pepin, K., Ellsworth, W. L., Zebkar, H., Segall, P., Gu, Y. J., and Samsonov, S., 2023, Disposal from in situ bitumen recovery induced the ML 5.6 Peace River earthquake: Geophysical Research Letters, **50**, No. 6, e2023GL102,940.

Uieda, L., Tian, D., Leong, W. J., Schlitzer, W., Grund, M., Jones, M., Fröhlich, Y., Toney, L., Yao, J., Magen, Y., Tong, J.-H., Materna, K., Belem, A., Newton, T., Anant, A., Ziebarth, M., Quinn, J., and Wessel, P., 2023, PyGMT: A Python interface for the Generic Mapping Tools.

Vasyura-Bathke, H., Dettmer, J., Biegel, K., Salvage, R. O., Eaton, D., Ackerley, N., Samsonov, S., and Dahm, T., 2023, Bayesian inference elucidates fault-system anatomy and resurgent earthquakes induced by continuing saltwater disposal: Nature Communications Earth & Environment, accepted.

Woollam, J., Münchmeyer, J., Tilmann, F., Rietbrock, A., Lange, D., Bornstein, T., Diehl, T., Giunchi, C., Haslinger, F., Jozinović, D., Michelini, A., Saul, J., and Soto, H., 2022, SeisBench—A Toolbox for Machine Learning in Seismology: Seismological Research Letters, **93**, No. 3, 1695–1709, https://pubs.geoscienceworld.org/ssa/srl/article-pdf/93/3/1695/5596371/srl-2021324.1.pdf.
URL https://doi.org/10.1785/0220210324

Woollam, J., Rietbrock, A., Bueno Rodriguez, A., and De Angelis, S., 2019, Convolutional neural network for seismic phase classification, performance demonstration over a local seismic network: Seismological Research Letters, **90**.

Wozniakowska, P., and Eaton, D. W., 2020, Machine Learning-Based Analysis of Geological Susceptibility to Induced Seismicity in the Montney Formation, Canada: Geophysical Research Letters, **47**, No. 22, e2020GL089,651, e2020GL089651 2020GL089651.

Wu, J., Chen, X.-Y., Zhang, H., Xiong, L.-D., Lei, H., and Deng, S.-H., 2019, Hyperparameter optimization for machine learning models based on bayesian optimizationb: Journal of Electronic Science and Technology, **17**, No. 1, 26–40.
URL https://www.sciencedirect.com/science/article/pii/S1674862X19300047

Zhu, W., and Beroza, G. C., 2018, PhaseNet: a deep-neural-network-based seismic arrival-time picking method: Geophysical Journal International, **216**, No. 1, 261–273.
URL https://doi.org/10.1093/gji/ggy423

Zhu, W., McBrearty, I. W., Mousavi, S. M., Ellsworth, W. L., and Beroza, G. C., 2022, Earthquake Phase Association Using a Bayesian Gaussian Mixture Model: Journal of Geophysical Research: Solid Earth, **127**, No. 5, e2021JB023,249.